

Remarks/Arguments

In the Non-final Office Action mailed on 25 July 2006, the Examiner rejected claims 1-5, 12-14, 17 and 18 under 35 U.S.C. §103(a) as unpatentable over Bakke (United States Patent Publication Number 2005/0071532) in view of Badamo (United States Patent Publication Number 2002/0181476) and further in view of Bakke (U.S. Patent No. 6,704,812), and rejected claims 6-9 and 15-16 as unpatentable over Bakke '532 in view of Bakke '812 in view of Badamo and further in view of Barrow (United States Patent Publication Number 2002/0188786).

Applicants respectfully traverse the Examiner's rejections under §103, and respectfully request reconsideration and withdrawal of the rejections.

§103 Claim Rejections

The Examiner rejected claims 1-5, 12-14, 17 and 18 under 35 U.S.C. §103(a) as unpatentable over Bakke '532 in view of Badamo and further in view of Bakke '812. The Examiner rejected claims 6-9 and 15-16 under 35 U.S.C. §103(a) as unpatentable over Bakke '532, Badamo and Bakke '812 in further view of Barrow. Though Applicants continue to take issue with the Examiner's suggested motivation for the proposed combinations, the Examiner's combinations still fail to teach or reasonably suggest every element of the rejected claims. The rejections will be discussed in regard to independent claim 1.

Independent claim 1 recites a multi-chip module (MCM), a well known term of art in the electronics design industry. For example, the Electronic Design Automation Consortium provides a glossary definition of MCM (multi-chip module) as follows:

A type of board technology whereby multiple, unpackaged integrated circuits (bare die) are mounted — along with signal conditioning or support circuitry such as capacitors and resistors — on a single laminate or ceramic base material. The MCM footprint is much smaller than conventional single chip packages, resulting in a smaller motherboard and smaller space requirements for panels, enclosures, and cabling. The result is a high-density module that resembles a single component when mounted on a printed circuit board. By combining more circuit functions in an MCM, fewer system assemblies are required, resulting in lower

circuit design costs, integrated functional testing, and higher manufacturing yields.

(see http://www.edac.org/industry_glossary.jsp)

The MCM comprises a plurality of SAS expander component circuits, i.e., a device for expanding the number of ports of a SAS network domain used to interconnect SAS initiators and SAS targets. Each SAS expander component circuit has a number of internal ports that are internal to the MCM. Each SAS component expander has a number of external ports for coupling to SAS devices (e.g., targets and initiators) external to the MCM. Selected SAS expander circuits are coupled through selected ones of their internal ports to an internal fabric during configuration of the coupling. The configuration determines paths between various internal ports of the SAS expander circuits. Once the MCM is initialized, the paths remain static until the next initialization (e.g., the next time the MCM is reset). The MCM further comprises coordination logic communicatively coupled to the SAS expander component circuits to coordinate operation between the plurality of SAS expander component circuits. The coordination logic allows the MCM to be presented as a single expander to devices outside the module (e.g., a SAS expander having a single SAS address that performs SMP exchanges as a single SAS address). Thus, the MCM allows for flexible extension and configuration of SAS expander ports.

REFERENCES FAIL TO SHOW A MULTI-CHIP MODULE (MCM)

First and foremost, nothing in the art of record (considered individually or in any combination) teaches or reasonably suggests a SAS expander implemented as a Multi-Chip Module ("MCM") - the very essence of the invention. None of the references describe an MCM as the term is understood by those of ordinary skill in the art. This term of art is well understood to describe a single integrated circuit that is manufactured to incorporate one or more other integrated circuit dies or components. Such MCM devices are generally fabricated using integrated circuit fabrication techniques - notably

distinct from fabrication/assembly of printed circuit board cards and/or entire systems incorporating numerous such cards.

The references teach nothing of such an MCM structure for a SAS expander but rather teach structures of printed circuit cards and systems incorporating a plurality of such cards. The Examiner stated in the Office Action that Bakke '532 teaches such an MCM structure. Specifically, the Examiner points to item 100 of Figure 4 of Bakke '532. The Examiner states that Bakke '532 teaches multiple chips (Figure 4, items 0, 1, 102) disposed onto a module (Figure 4, item 100). Item 100 of Figure 4 of Bakke '532 is a SAS network comprising a pair of edge expanders and a plurality of SAS initiators and devices (paragraph 17 of Bakke '532). In other words, item 100 is an entire SAS system domain coupling multiple initiators to multiple targets through multiple paths via intervening expanders.

As the Examiner correctly points out, Bakke '532 does not disclose the components of Figure 4 incorporated into a single integrated circuit that is manufactured to incorporate one or more other integrated circuit dies or components. The Examiner states that it would have been an obvious as a matter of design choice to incorporate edge expanders 0, 1, 102 of Bakke '532 into a single integrated circuit, since such a modification would have involved a mere change in size of a component.

The Applicants respectfully disagree with the Examiner's assertion that it would have been obvious to integrate both edge expanders of Bakke '532 into a single integrated circuit. Such integration of the edge expanders of Bakke '532 into a single integrated circuit would be more than a mere change in the size of a component. Such integration would fundamentally change the teachings of the Bakke '532 reference.

Amended claim 1 is not mere integration of components that is a change in size of the components. Rather, the integration of components offers features not found in the prior art. Specifically, the MCM offers multiple SAS expander component circuits presented as a single (i.e., unified) expander (e.g., a SAS expander with a single address). Such a single expander is not possible with the external, separated and standalone

expanders of Bakke '532 or any of the prior art or record, considered individually, or in any combination.

Bakke '532 specifically teaches away from integrating edge expanders 0, 1, 102 into a single integrated circuit. The object of the Bakke' 532 reference is to provide a method and apparatus for implementing resilient connectivity in a Serial Attached SCSI (SAS) domain (Abstract of Bakke '532). Bakke '532 discusses how prior SAS networks suffer from connectivity problems when a component (e.g., an edge expander) fails, and the path between an initiator and target is lost. Bakke '532 further discusses how redundancy techniques in SAS networks have failed in the past due to the failure of one or more components or ports of the SAS network.

Bakke '532 solves connectivity problems by providing multiple independent and distinct edge expanders (paragraphs 0013 and 0021 of Bakke '532). Two redundant edge expanders are provided within the SAS network 100 of Bakke '532 to provide redundant pathways for requests should one or more components or ports within the SAS network fail. Thus, if one edge expander fails, then communications may be routed between the initiators and devices through another edge expander. Or, if one port of the multi-ported devices fails, data may none the less be routed through another port of the device coupled through a second expander.

Such redundant and resilient benefits are lost if both edge expanders are integrated into a single integrated circuit as recited in rejected claim 1. As would be typical for any integrated circuit design, the entire circuit typically fails for all practical purposes. In the context of the MCM of claim as potentially read on the structure of Bakke '532, in any typical failure of the MCM, both edge expanders would become unavailable to SAS devices coupled to the edge expanders. Any benefits derived from redundancy as described by Bakke '532 would be lost, as the pathways between initiators and targets would be unavailable. Further, redundancy benefits at the heart of the Bakke teachings would be negated by incorporating both of the edge expanders of Bakke '532 onto a single integrated circuit (not to mention also integrating the initiators 1-2 and devices 1-4 of SAS network 100 onto the same single integrated MCM circuit). If

Bakke's edge expanders integrated on such a single integrated circuit both fail, then the devices and initiators would no longer be useful, as the pathways between the devices and initiators would be unavailable.

Thus, Applicants submit that Bakke '532 does not teach an MCM, integrated circuit component. Further, there is no suggestion or motivation in the art to produce or design a SAS expander as a single, MCM, integrated circuit component found in Bakke '532, or any of the other cited art, considered individually, or in any combination.

REFERENCES FAIL TO SHOW COORDINATION LOGIC PRESENTING A SINGLE EXPANDER

Bakke '532 additionally fails to disclose coordination logic as recited by claim 1. The coordination logic of claim 1 coordinates operation of the SAS expander component circuits to present a single expander (e.g., a SAS expander with a single SAS address) to devices outside the MCM. The Examiner states that page 2 paragraph 22 of Bakke '532 discloses coordination logic communicatively coupled to the plurality of SAS expander component circuits to coordinate operation of the plurality of SAS expander component circuits. Specifically, the Examiner states that paragraph 22 describes "receipt of data from one of the devices cause the edge expanders to use logic to determine where the data is to be sent, therefore, it would be obvious to one of ordinary skill in the art that there is coordination logic within the edge expanders." Paragraph 22 of Bakke '532 discusses the routing of edge expanders 0-1. When a frame is received at a particular direct routing port, the edge expander compares the destination SAS address contained within the frame to the SAS address of each of the nodes of the other direct routing ports of SAS devices coupled to the edge expander. If a matching address is located, then the frame is routed to the node. Otherwise, the frame is sent to subtractive routing port 120 for routing to another edge expander.

A subtractive routing port is used when a SAS node makes a connection request to an edge expander requesting a node that is not directly connected to the edge expander

(paragraph 0008 of Bakke '532). The edge expander does not know how to route requests for end devices which are not directly connected to the edge expander (e.g., a request for a SAS device coupled to another edge expander). A request for a device not directly connected to the edge expander is forwarded out the subtractive routing port normally to a fanout expander, which contains a routing table. The edge expander is able to determine the correct edge expander to route the request (e.g., the edge expander to which the requested node is directly connected) (paragraph 0009 of Bakke '532). The request is forwarded by the fanout expander to the correct edge expander to which the requested node is directly connected.

As noted in Bakke '532, the edge expander is a simple device with significantly less function than a switch would have (paragraph 0008 of Bakke '532). The edge expanders have no routing tables (paragraph 0008 of Bakke '532). If a requested SAS address isn't found directly coupled to the edge expanders of Figure 4 of Bakke '532, then the request is forwarded out the subtractive routing port to the other connected edge expander. The other edge expander then checks to see if the requested SAS address is connected to one of its own direct ports. If not, the frame is rejected. If the edge expander had coordination logic, then it would be able to correctly route a request to the correct destination edge expander. However, the edge expander doesn't have coordination logic. If an edge expander doesn't know where to route a request, then the request is just forwarded to the next link through the subtractive routing port in hope that the next edge expander will know where to route the request. The SAS network 100 of Bakke '532 eliminates the fanout expander, and thus, the routing logic of SAS network 100.

Assuming, *arguendo*, that the edge expanders comprise coordination logic, then Bakke '532 still does not disclose that the coordination logic is adapted to present a single expander to devices outside the module, wherein the single expander performs SCSI management protocol ("SMP") exchanges as a single SAS address as recited by amended claim 1. Both edge expanders in Figure 4 of Bakke '532 are totally independent, each has different addresses, and each perform separately in SMP exchanges. Thus, the pair of edge expanders doesn't present a single expander to devices outside the module which

respond to SMP requests as a single address. The Examiner correctly notes that the edge expanders of Bakke '532 do not present a single expander to devices outside the module. The Examiner does not recognize that the important functionality of the coordination logic, i.e., presenting a single expander to devices outside the module, is not shown in the Bakke '532 reference. The Examiner refers to another reference to find this recitation, which is important aspect of the coordination logic of the MCM.

The devices and initiators connected to the edge expanders in Bakke '532 are configured as connected to two different edge expanders. Communications can be routed by the initiators and devices through either of two different ports of the device, each connected to a different edge expander. SMP exchanges for each edge expander in Bakke '532 are addressed separately to each of the expanders. For example, discovery related exchanges as the SAS domain is initialized involve each of the multiple expanders in Bakke '532 individually, and not as a single expander. Each SMP request must be sent to a different SAS address for each of the edge expanders. Thus, the edge expanders are presented as two edge expanders having two different SAS addresses. Each edge expander responds to SMP requests as a different device at a different SAS address.

Admitting that Bakke '532 fails to show such a single presentation of the multiple expander components as claimed, the Examiner states that Bakke '812 presents a single expander (e.g., a unified expander) to devices outside a module. Bakke '812 is related to management of redundant physical paths by a host computer system to peripheral devices. The Examiner cites to Bakke '812 column 4, lines 52-57, "showing a plurality of paths connecting to various components within the module are shown as one path [and therefore, one component to the operating system]". The cited passage generally describes that a redundancy manager (i.e., within a host adapter) may comprise a path resolver that resolves all the independent physical paths to the peripheral device into one logical path presented to the operating system.

While the operating system is presented with a single pathway, the host system is still presented with multiple pathways, which may be considered multiple devices. From the host system's perspective, there are still multiple devices and/or multiple pathways attached to the host system, and the redundancy manager is a layer within the host system

that handles the actual communication to the devices so that the operating system does not need to worry about handling multiple protocols and pathways. However, each device and pathway needs to have multiple addresses for routing and configuration. Otherwise, the redundancy manager can't configure or reach the device. Thus, Bakke '812 does not disclose presenting a single expander to devices outside the module. Further, neither of the Bakke references, nor any of the art of record, considered individually, or in any combination, discloses multiple expander circuits perform SMP exchanges (e.g., a configuration request) as a single address. Thus, Applicants submit that none of the references, considered individually, or in any combination, disclose the coordination logic as recited by amended claim 1.

REFERENCES FAIL TO SHOW CONFIGURATION OF COUPLING OF THE
INTERNAL PORTS AS STATIC FOLLOWING INITIALIZATION

The internal fabric of claim 1 provides static connections/routes among the internal ports of the various SAS expander component circuits of the MCM. The static connections/routes are static (as discussed above) in that they remain unchanged during operation of the MCM. The configuration may be determined either at time of manufacture of the MCM or based on programmable logic used when the MCM is reset (e.g., powered up or otherwise reset to a powered up state). Nothing in the prior art of record, alone or in any combination, shows such a SAS expander system/circuit design in which the couplings/routes within the MCM implemented SAS expander, once determined, remain static for the remaining period of operation of the MCM only to be changed upon a reset of the MCM (if ever changed after manufacture).

The Examiner suggests that Badamo provides such a teaching and points to Badamo element 20 of Figure 3 and paragraphs 41 and 43 in support thereof. Even assuming Badamo (or any of the applied art) is properly combined with the other applied references, nothing in Badamo teaches that the "fabric card" (FC 20) of Figure 3 provides a static routing structure following initialization of the MCM SAS expander. To the contrary, Badamo teaches that the routing through FC 20 may be altered dynamically to

allow for redundant backup of one LC 22 or SC 24 by another LC 22 or SC 24. Thus FC 20 is anything but static!

Rather, Badamo teaches that a control card (CC 36) of Figure 3 configures the FC 20 and may reconfigure the dynamic FC 20 as required during operation of the system. As stated in Badamo, LC 22 with the capability to classify ingress traffic can thus make use of unused capacity on the ingress service cards 24 by changing the routing. Even though Badamo uses the word static, the disclosed card is clearly programmable during operation of the system and not static as claimed in the rejected claims. Thus, Applicants submit that none of the references, considered individually, or in any combination, disclose coupling of the internal ports as static following initialization of the MCM, as recited by amended claim 1.

In view of the above discussion, Applicants assert that independent claim 1 is allowable at least for the reasons stated above. At least these same arguments apply to independent claims 17 and 18, rejected by the Examiner for similar reasons and to dependent claims 2-9 and 12-14. Applicants respectfully request reconsideration and withdrawal of the outstanding rejections.

Conclusion

Applicants have amended claims 1 and 17-18 for editorial clarity and to better protect the invention. Applicants have addressed each issue raised by the Examiner and respectfully request reconsideration and withdrawal of all outstanding rejections and passage of the application to allowance and issue.

Applicants believe no fees are due in this matter. Should any issues remain, the Examiner is encouraged to telephone the undersigned attorney.

Respectfully submitted,



Max S. Gratton (#56,541)
Duft Bornsen & Fishman, LLP
1526 Spruce Street, Suite 302
Boulder, CO 80302
(303) 786-7687
(303) 786-7691 (fax)